Learning from Inflation Experiences*

Ulrike Malmendier[†]
UC Berkeley, NBER and CEPR

Stefan Nagel[‡]
Stanford University and NBER

February 2009

Abstract

We investigate how subjective expectations about future inflation rates are shaped by individuals' life-time experiences of inflation. Using more than 50 years of microdata on inflation expectations from the Reuters/Michigan Survey of Consumers, we find that differences in experiences correlate with differences in subjective expectations. In particular, young indvidiuals place more weight on recently experienced inflation than older individuals. As a result, periods of high surprise inflation like in the 1970s lead to substantial disagreement between young and old individuals about future inflation rates. Our results can be interpreted as adaptive learning, but in a way that indviduals learn from data experienced over their life-times, rather than from all "available" historical data. An aggregate measure that summarizes the average life-time inflation experiences of individuals at a given point in time is useful in predicting excess returns on long-term bonds.

*We thank Ralph Koijen, Monika Piazzesi and workshop participants at Stanford and the American Economic Association meetings for helpful comments and Sith Chaisurote for excellent research assistance.

[†]Department of Economics, University of California, Berkeley, ulrike@econ.berkeley.edu

[‡]Stanford University, Graduate School of Business, 518 Memorial Way, Stanford, CA 94305, e-mail: Nagel_Stefan@gsb.stanford.edu

1 Introduction

That macroeconomic outcomes and asset prices depend in crucial ways on expectations of economic actors is well understood at least since Keynes (1936). But we know little about how economic agents form their subjective beliefs about the future. The literature on boundedly rational adaptive learning (see Bray (1982); Sargent (1993); Evans and Honkapohja (2001)) studies the implications of individuals acting as econometricians and making forecasts based on historical data, but there is yet little direct empirical evidence on the actual forecasting rules employed by firms, individuals, and policy makers, even though understanding the formation of inflation expectations, and macroeconomic expectations more generally, is likely to be of first-order importance for macroeconomic policy (Bernanke (2007)).

We examine one particular expectations formation mechanism in this paper, namely that individuals' expectations are particularly strongly influenced by their own "experiences", by which we mean macroeconomic data that individuals experienced during their life time. This *learning-from-experience* hypothesis is related to the adaptive learning approach in macroeconomics, but it differs in a key respect. Suppose that two individuals (A and B) perceive a macroeconomic variable, say stock market returns, to be *iid*. According to the least-squares learning rules popular in the adaptive learning literature, these individuals would form their subjective expectations about next year's return as a simple average of past realized returns, using all available historical data (whatever "all available" might actually mean in practice). In contrast, our learning-from-experience hypothesis posits that individuals are more strongly influenced by data realized during their life-times than by other historical data. In this simple example, this would lead to several important implications: First, there would be heterogeneity in beliefs. Individuals who lived through periods of higher stock market returns would have more optimistic beliefs about future returns than individuals who experienced mostly poor returns during their life times. Second, learning dynamics are perpetual. Even in this simple *iid* setting, beliefs would keep fluctuating and would never converge in the long-run, as memories of historical data get lost as new generations emerge. Third, it should be possible to estimate the learning-from-experience effect from cross-sectional differences in subjective expectations between individuals of different age. This would make it possible to control at the same time for other unobserved macroeconomic influences on expectations.

Our analysis is motivated by earlier empirical findings of Malmendier and Nagel (2008), from data from the Survey of Consumer Finances (*SCF*), that various measures of individuals' risk-taking and portfolio allocations are correlated with individuals' macroeconomic experiences. Their data, however, did not allow to determine whether these effects are driven by beliefs (e.g., experiences of high stock returns make individuals more optimistic) or by endogenous preferences (e.g., experiences of high stock returns make individuals less risk averse or lead to other changes in "tastes" for certain asset classes). In this paper, we use direct data on expectations to focus specifically on the beliefs channel. To this end, we focus on subjective inflation expectations, for which we have household-level microdata from the Reuters/Michigan Survey of Consumers (*MSC*) spanning a period of more than 50 years. We use this inflation expectations data to estimate how subjective beliefs about future inflation are related to individuals' life-time experiences of inflation.

We implement the estimation of experience effects by estimating a learning rule similar to Marcet and Sargent (1989), but with the twist that we allow individuals to learn only from data realized during their life-time. Individuals forecast inflation by forming a weighted average of inflation rates experienced in the past. The learning rule features a decreasing "gain", which means that adding an additional inflation observation has less of an impact for older individuals who have more data accumulated in their life-time histories. A gain parameter determines how fast these gains decrease as more data accumulates. We estimate this gain parameter by fitting the learning rule to individuals' reported inflation expectations in the *MSC* (with forecast horizons of 1 and 5-10 years, respectively). The estimates of this gain parameter then effectively tells us how people weight their inflation experiences to come up with their current beliefs.

The availability of microdata is crucial for our purpose, as it allows us to properly identify the experience effect. First of all, the learning-from-experience hypothesis predict heterogeneity in subjective beliefs, and hence only microdata allows to uncover whether this predicted heterogeneity exists. Moreover, the cross-sectional dimension of microdata allows us to employ time dummies in the estimation, which makes it possible separate experience effects from other observed and unobserved influences on expectations. For example, we do not need to assume that past inflation experiences are the only influence on people's subjective beliefs about future inflation. The time dummies absorb any other variation in inflation expectations that is common to all individuals, which could, for example, include time-varying beliefs in the credibility of monetary policy makers in fighting inflation, or some reliance on published forecasts by professional forecasters. The presence of time dummies also implies that we do not implicitly assume that data realized prior to an individuals' life-time is completely ignored. Instead, the component of expectations that our method attributes to learning-from-experience reflects the incremental explanatory power that life-time experiences have in explaining inflation expectations, over and above any common influences that individuals may be subject to and that are absorbed by the time dummies.

Put differently, our identification strategy relies on time-variation of cross-sectional differences in inflation expectations, and looks at how these are related to time-variation in cross-sectional differences of inflation experiences. This is a big advantage of the learning-from-experience setting, compared with earlier approaches in the literature, where researchers have estimated adaptive learning models or other models of belief formation by fitting them to aggregate time-series of expectations (e.g., median expectations). With aggregate data, it is difficult to rule out that correlated unobserved omitted macroeconomic variables could bias the estimation results. With our approach such unobservables are absorbed by the time dummies.

Our estimation results show that learning from experience has an economically important effect on inflation expectations. Individuals of different ages often differ substantially in their inflation expectations, and these differences are related to heterogeneity in their experienced inflation histories. These differences are especially large in times when recent inflation rates were far from the historical mean. For example, in the late 1970s and early 1980s, the average inflation expectations at a 5-10 year horizon of individuals under the age of 40 exceeded those of older individuals above age 60 by several percent, consistent with the fact that the life-time experience of younger individuals was dominated by the high inflation years of the 1970s, while the experience of older individuals also included the low inflation years in the 1950s and 1960s. This difference in expectations between young and old only slowly faded away until the mid-1990s after a many years of moderate inflation.

The estimates of the gain parameter reveal how much weight people put on recently experienced inflation rates relative to inflation experienced earlier in life. Our estimates imply that recent inflation experiences receive relatively higher weight, but for older individuals experiences from 20 to 30 years ago can still have some long-run effect.

In our estimation, we also control for differentials in inflation rates faced by the young and the elderly, using the experimental CPI for the elderly, avaible since the early 1980s from the Bureau of Labor statistics. While there are persistent differences between the overall inflation rate and the inflation rate of the elderly, these inflation differentials do not help to explain our results.

We also investigate the implications of the estimated learning-from-experience effects for the economy in aggregate. Using our estimates of how individuals of different ages weight their past inflation experiences, we calculate the learning-from-experience component of inflation expectations of each age group. We then form a wealth-weighted average across all age groups. This aggregate experienced inflation measure summarizes the average inflation experience of economic agents in a given quarter and we find that it predicts aggregate forecast errors, as well as the excess returns on long-term government bonds. In times when the average individual experienced high past inflation, the learning-from-experience component pushes up subjective inflation expectations, and it only slowly decays in response to declining inflation, giving rise to persistent forecast errors. At the same time, it apparently also leads to depressed long-term bond prices, resulting high excess returns on long-term bonds in the future. These results underscore that learning-from-experience effects have potential to explain macroeconomic phenomena.

Our paper connects to a number of works in the literature. Conceptually, our approach is related to Honkapohja and Mitra (2003) who consider a model of bounded memory learning. Bounded memory learning is similar to the learning-from-experience framework that we use here in that memory of past data is lost, but in bounded memory learning agents are homogeneous, whereas in our setting agents' memory depends on their age.

There is only a small, but growing literature that looks at heterogeneity in expectations formation with micro data. Building on early work by Cukierman and Wachtel (1979), Mankiw and Wolfers (2003) examine the time-variation in dispersion in inflation expectations, and they relate it to models of "sticky" information.Carroll (2003) further investigates the sticky information model, but with aggregate data on inflation expectations. Branch (2004) and Branch (2007) estimate from survey data how individuals choose among competing forecasting models.

Orphanides and Williams (2006) and Milani (2007) set up macroeconomic models in which agents learn adaptively, which helps explain inflation persistence and inertia in other macroeconomic variables. But the learning dynamics are specified ad-hoc, in Orphanides and Williams' case, or estimated by fitting macroeconomic data, not expectations data, in Milani's case, and there is no allowance for expectations heterogeneity. Piazzesi and Schneider (2008a) incorporate data survey data on heterogeneous subjective inflation expectation in asset pricing, Piazzesi and Schneider (2008b) use data on subjective interest rate expectations and a model with adaptive learning. Our paper provides an alternative and complementary perspective on expectations heterogeneity via the learning-from-experience channel.

Some suggestive evidence for the existence of learning-from-experience effects is presented in Greenwood and Nagel (2009) and Vissing-Jorgensen (2003), who show that young mutual fund managers and young individual investors in the late 90s were more optimistic about stocks, and in particular technology stocks, than older investors, consistent with young investors being more strongly influenced by their recent good experience with technology stocks. Kaustia and Knüpfer (2008) find that investors' personal experience with returns from investing in initial public offerings affects their willingness to subscribe to future offerings. Vissing-Jorgensen (2003) also points out that there is age-heterogeneity of inflation expectations in the late 1970s and early 1980s. Our paper investigates learning-from-experience effects in a more systematic fashion, with a long-term data set, and explicit estimation of a learning-from-experience forecasting rule.

The rest of the paper is organized as follows. Section 2 introduces our analytic framework on learning from experience and our estimation approach. Section 3 discusses the data set on inflation expectations. Section 4 presents our core set of results on learning-from-experience effects in inflation expectations. In Section 5 we look at the implications of our results for forecast errors and in Section 6 we show how learning-from-experience effects can help understand predictable variation in excess returns of long-term bonds. Section 7 concludes with some final thoughts.

2 Learning from experience

Consider a two individuals, one is member of the cohort born at time s, and the other belongs to the cohort born at time $s+j$. Suppose that at time t they form expectations of next period's inflation, π_{t+1}, based on the history of past inflation rates. The essence of the learning-from-experience hypothesis is that when these two individuals forecast π_{t+1}, they draw on inflation histories of different lengths, and they place a different weights on recent and distant historical data. As a result, two individuals of different cohorts may produce different forecasts at the same point in time. Our goal is to investigate whether individuals' inflation forecasts are indeed influenced in such differential ways by historical inflation rates.

To set up an analytical framework, we need to have some prior idea as to how individuals' forecasting rules might look like. Our approach is guided by the prior literature in two ways. First, the candidate forecasting rule we examine has close resemblance to those in the adaptive learning literature, in particular Marcet and Sargent (1989) (see also Sargent (1993) and Evans and Honkapohja (2001)). The key departure from standard adaptive learning models is that we allow individuals to put more weight on data experienced during their lifetimes than on other historical data, which results in cross-sectional heterogeneity in expectations between members of different cohorts. Second, our candidate forecasting rule sets expectations equal to a weighted average of past inflation rates, motivated by the evidence in Malmendier and Nagel (2008) that individuals' portfolio allocations to risky assets are positively correlated with a weighted average of risky asset returns realized during individuals life-times. Thus, in this paper we ask whether such a weighted average of experienced macroeconomic data can also explain subjective inflation expectations data.

That individuals set inflation expectations equal to a weighted average of past inflation rates would arise in standard adaptive learning framework if individuals' perceived law of motion of inflation is a stochastic trend model of the sort considered by Nelson and Schwert (1977), Barsky (1987),

$$\pi_t = \tau_t + \eta_t. \tag{1}$$

$$\tau_t = \tau_{t-1} + \varepsilon_t. \tag{2}$$

according to which individuals believe that the trend in inflation, τ_t, follows a random walk.[1] With recursive least-squares learning (see Evans and Honkapohja (2001)), individuals' time-t estimate of τ_{t+1}, denoted with $\hat{\tau}_{t+1|t}$, is obtained from

$$\hat{\tau}_{t+1|t} = \hat{\tau}_{t|t-1} + \gamma_t \left(\pi_t - \hat{\tau}_{t|t-1} \right), \tag{3}$$

where the recursion is started at some point in the (distant) past. The gain parameter γ_t determines the degree of updating when faced with the inflation surprise $\pi_t - \hat{\tau}_{t|t-1}$. With the gain sequence $\gamma_t = 1/t$, this algorithm would simply produce the sample mean of past inflation, using all data available until time t (assuming initial values are set appropriately). With γ_t set to a constant, the algorithm becomes a constant-gain learning algorithm. Constant-gain learning implies that past data is weighted with geometrically decaying weights.[2] Thus, in this adaptive learning setting, individuals would use some weighted-average of past inflation rates to estimate the trend in inflation.

Such relatively simple learning algorithms are motivated in the adaptive learning literature by the fact that economic agents face cognitive and computational constraints that limit their ability to use optimal forecasts. Recursive least-squares algorithms are viewed as an approximation the "rules of thumb" that practitioners and individuals might employ. The focus of much of the adaptive learning literature is on the conditions under which such simple learning rules can lead to convergence to rational expectations. Our objective is different. We use this simple recursive least-squares learning algorithm as a starting point for an empirical investigation of individuals' actual forecasting rules, and we depart from the standard adaptive learning framework in some ways to allow for learning-from-experience effects.

Our key modification of the standard recursive least-squares learning framework is that we let the gain parameter depend on the age $t-s$ of the members of the cohort s, instead of the calendar time t,

$$\hat{\tau}_{t+1|t,s} = \hat{\tau}_{t|t-1,s} + \gamma_{t-s} \left(\pi_t - \hat{\tau}_{t|t-1,s} \right), \qquad t \geq s. \tag{4}$$

[1] According to Stock and Watson (2003), a heteroskedastic version of this model, where the variances of η_t and ε_t are time-varying is an empirically plausible description on the process of U.S. inflation. We abstract from the implications of such heteroskedasticity, as our focus is not on deriving an optimal forecasting rule, but rather to make a plausible guess about an approximation to the simple forecasting rules that individuals might employ, in order to have a starting point for our empirical investigation.

[2] For the stochastic trend model (1) with homoskedastic residuals constant gain is optimal as $t \to \infty$, and its magnitude depends on the relative variance of η_t and ε_t. With a finite sample of data, the optimal gain declines initially (Kalman (1960); Muth (1960)).

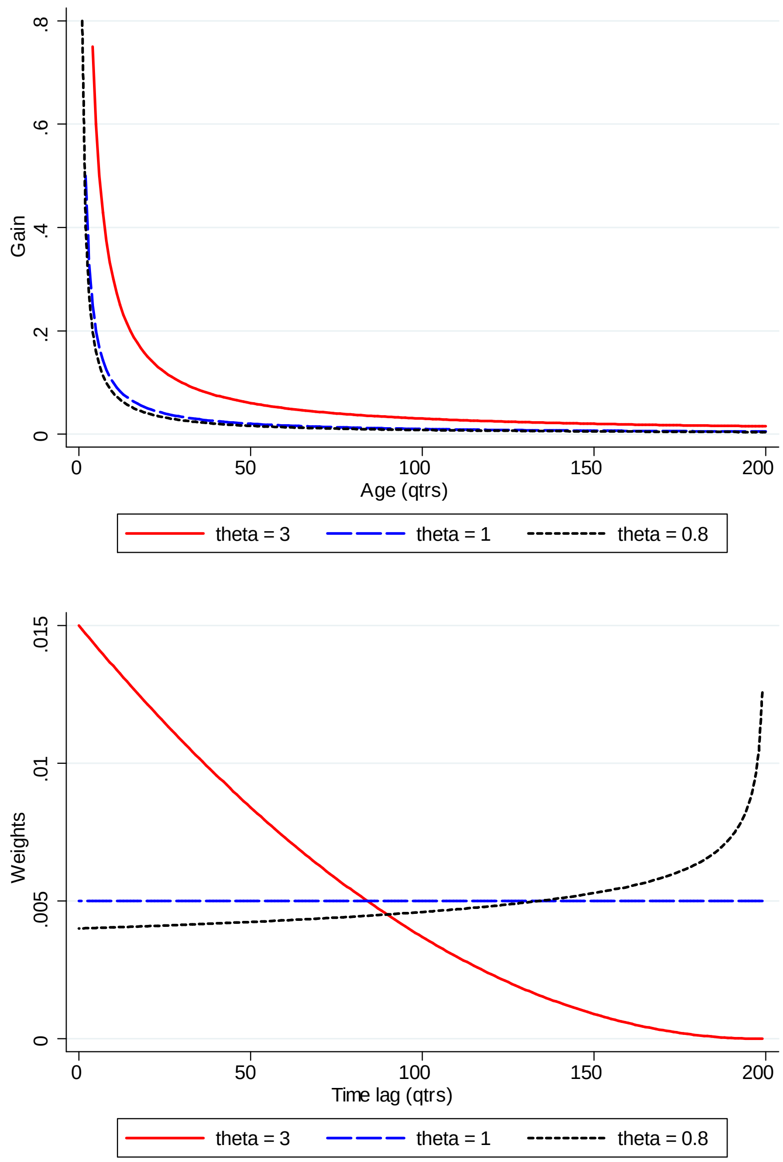

Figure 1: Examples of gain sequences and associated weights of past inflation rates

As a result, individuals of different age can be heterogeneous in their forecasts $\hat{\tau}_{s,t}$ and they adjust their forecasts to different degrees in response to surprise inflation. Moreover, we let the recursion start with $\gamma_1 \geq 1$ and $\hat{\tau}_{s+1|s,s} = \pi_s$, which implies that data before birth is ignored (but, as will become clear below, our econometric specification does not assume that data before birth is necessarily ignored).

We consider the following decreasing-gain specification,

$$\gamma_{t-s} = \frac{\theta}{t-s}, \tag{5}$$

where θ is a constant parameter that determines the shape of the implied function of weights on past experienced inflation observations. This gain sequence resembles those used, for example, in Marcet and Sargent (1989), but with the difference that the gain here is decreasing in age, not in time, and individuals use only data realized during their life-times, as opposed to all historical data. That the gain decreases with age is a sensible assumption in context of the learning-from-experience hypothesis. Young individuals, who have experienced only a small set of historical data, should presumably have a higher gain than older individuals, who have experienced a longer data history, and for whom a given time-t observation should have less marginal impact on their estimates.

To illustrate the role of the parameter θ, Figure 1 presents the sequences of gains (in the top graph) as a function of the age of the individual and the implied weights (in the bottom graph) on past inflation observations as a function of the time lag relative to current time t. In this example, we consider an individual with 50 years of age (200 quarters). Focusing on the top graph first, it is apparent that gains decrease with age at a slower rate when θ is higher. This means that less weight is given to observations that are more distant in the past, as shown in the bottom graph. For $\theta = 1$, for example, all historical observations since birth are weighted equally. For $\theta > 1$ weights on earlier observations are lower than those on more recent observations. The figure shows the example for $\theta = 3$, where very little weight is put on observations in the first 50 quarters since birth towards the right of the bottom graph.

We show in Appendix A.3 that the decreasing-gain learning scheme in Eq. (5) produces weight sequences that are virtually identical to those produced by the weighting scheme in Malmendier and Nagel (2008) for appropriate choices of the weighting parameters. This allows us to easily compare the implied weights obtained by estimating the parameter θ from inflation expectations data, with the earlier evidence in Malmendier and Nagel where the weighting scheme is estimated from data on portfolio allocations.

It would not be realistic to assume that individuals' expectations formation is influenced *only* by past inflation data and *only* by data realized during their life-times, and so in setting up our econometric specification we consider the more general model

$$\pi^e_{t+1|t,s} = \beta \hat{\tau}_{t+1|t,s}(\theta) + (1-\beta)\, x_t, \tag{6}$$

where $\pi^e_{t+1|t,s}$ is the time-t forecast of period $t+1$ inflation made by members of cohort s. This subjective expectation is a weighted average of the learning-from-experience component $\hat{\tau}_{t+1|t,s}(\theta)$ and an unobserved common component of individuals' forecasts x_t.

This unobserved component x_t could represent any kind of forecast based on common information available to all individuals at time t. For example, individuals might rely, to some extent, on the opinion of professional forecasters, and the representation of their opinions in the news media (e.g., as in Carroll (2003)). The learning-from-experience effect might only make an incremental contribution over and above the influence of professional forecasters on individuals' opinions, and so x_t could capture the common component of individuals' inflation expectations that is driven by the influence of professional forecasts. Alternatively, x_t could capture the common component of individual forecasts that is based on all the available historical data (whatever that might be in practice), as opposed to their life-time experiences. Again, β then only captures the incremental contribution of life-time experiences $\hat{\tau}_{t+1|t,s}(\theta)$ to $\pi^e_{t+1|t,s}$ over and above this common component. Thus, we do not assume that individuals *only* use data realized during their life-times, but our goal is to empirically isolate the incremental effect of life-time experiences of inflation on expectations formation.

We estimate the following modification of Eq. (6):

$$\tilde{\pi}^e_{t+1|t,s} = \beta \hat{\tau}_{t+1|t,s}(\theta) + \delta' D_t + \varepsilon_{t,s}, \tag{7}$$

where $\tilde{\pi}^e_{t+1|t,s}$ denotes measured inflation expectations from survey data. In this estimating equation, we absorb the unobserved x_t with a vector of time dummies D_t. We also add the disturbance $\varepsilon_{t,s}$, which we assume to be uncorrelated with $\hat{\tau}_{t+1|t,s}$, but which is allowed to be correlated over time within cohorts. It captures for example, measurement error in the survey data and idiosyncratic factors influencing expectations beyond those explicitly considered here. We use this specification to jointly estimate θ and β with non-linear least squares.

The presence of time dummies in Eq. (7) implies that we identify β and θ, and hence the learning-from-experience effect on expectations, from cross-sectional differences between the subjective inflation expectations of individuals of different age, and from the evolution of those cross-sectional differences over time. This is a major advantage over prior work which has estimated adaptive learning rules from aggregate data, e.g., time-series of median inflation expectations of the entire sample. If one finds a time-series relationship between time-t median inflation expectations expectations and lagged inflation rates, it is difficult to rule out that this correlation could be caused by some omitted and perhaps unobservable contemporaneous macroeconomic variable. In our setting, the time dummies absorb any such macroeconomic variables. In this way, we can control for many alternative hypotheses about how individuals might form their expectations of future inflation.

3 Data

To estimate the learning-from-experience model, we use long-term historical data on the consumer price index (CPI) originally used in Shiller (2005), and available (updated until the end of 2007) on Robert Shiller's website. Based on this series, we calculate quarterly observations of annual inflation rates as the percentage change in the CPI from end of quarter $t-4$ to end of quarter t. To illustrate the long-run variation in inflation rates, Figure 2 shows five-year moving averages of this inflation rate series.

The central part of our analysis is to relate this history of inflation rates to individuals' expectations about future inflation. To measure individuals' inflation expectations, we use microdata from the Reuters/Michigan Survey of Consumers (MSC), conducted by the Survey Research Center at the University of Michigan. These surveys were administered since the early 1950s, initially three times per year, then quarterly from 1960 through 1977, and monthly since 1978 (see Curtin (1982)). Several questions in these surveys are the basis for the calculation of the University of Michigan Consumer

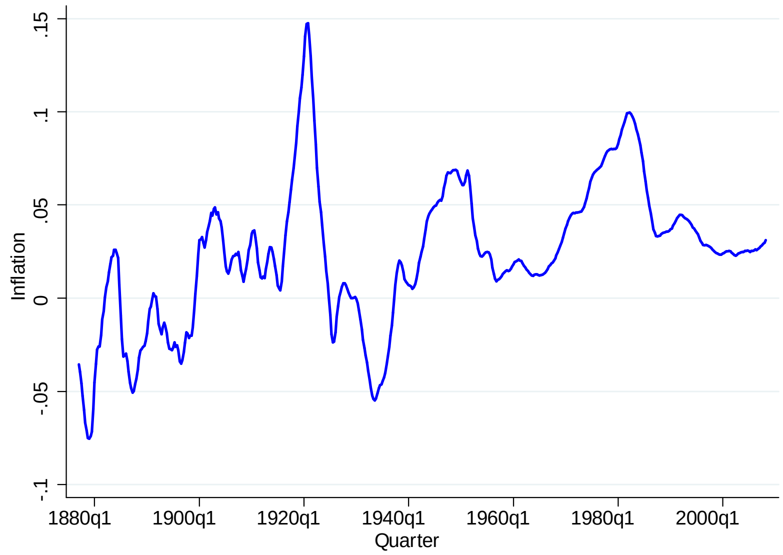

Figure 2: Five-year moving average of annualized CPI inflation rates

Sentiment Index.

We obtain surveys conducted during 1953 to 1977 from the Inter-university Consortium for Political and Social Research (*ICPSR*) at the University of Michigan. From 1959 to 1971, the questions of the winter-quarter Survey of Consumer Attitudes were administered as part of the Survey of Consumer Finances (*SCF*), and so we obtain those data from the *SCF* files, also available from the *ICPSR*. The data from 1978 to 2007 is available in from the University of Michigan Survey Research Center.

In most periods, survey respondents are asked two questions about expected inflation. One about the direction of expected future price changes ("up", "same", or "down") and one about the expected percentage change in prices. Moreover, in many periods, consumers are asked these two questions for both their expectations about price changes at a one-year horizon and over a 5-10 year horizon.

In our analysis, we focus on percentage expectations about future inflation. Figure 3 highlights the periods in which we have percentage expectations data at a one-year horizon (top graph) and 5-10 year horizon (bottom) graph. The quarters in which percentage expectations data is directly available in the survey data set are shaded in light grey. Those shaded in dark grey are quarters in which respondents in the survey are asked only the categorical question ("up", "same", or "down"). In those quarters we impute percentage responses from the categorical responses. The imputation procedure is described in detail in Appendix A.2.

Since our learning-from-experience hypothesis predicts that inflation expectations should be heterogeneous across different age groups, we aggregate the data at the cohort level. For each cohort defined by birth year, we compute, every quarter, the interpolated median of inflation expectations of members of this cohort.[3] If multiple surveys are administered within the same quarter, the interpolated median is taken across all surveys conducted within that quarter. We restrict our sample to respondents whose age ranges from 25 to 74. This means that for each cohort we obtain a quarterly series of inflation expectations that covers the time during which members of this cohort are from 25 to 74 years old.

To provide some sense of the variation in the data, Figure 3 plots the average inflation expectations of young (averaging across all cohorts that are in the age range from 25 to 39) and old (averaging across cohorts that are in the age range 61 to 75), relative to the interpolated median taken across all age groups. Thus, the figure plots cross-sectional differences, which we focus on in the estimation, not the time-variation in aggregate. For the one-year expectations in the top graph, the dispersion across age groups widens to 1-2% during the high inflation years of the 1970s and early 1980s. A much bigger dispersion is evident for the 5-10 year expectations in the bottom graph. Particularly during the time when inflation started to level off around 1982, the gap between young and old becomes very wide, reaching 3-4%. The fact that young individuals at the time expected higher inflation is consistent with the learning-from-experience story. The experience of young individuals at the time was dominated by the recent high-inflation years, while older individuals' experience also included the modest inflation rates of earlier decades.

[3] Percentage responses in the *MSC* are given in full (integer) percentage points. After ranking the N observations in the data, the interpolated median forms a weighted average of the next lowest reponse value below the median and the next highest respons value above the standard median, where the weights reflect the relative number of ranked observations with response equal to the median that are above or below the $N/2$ mark.

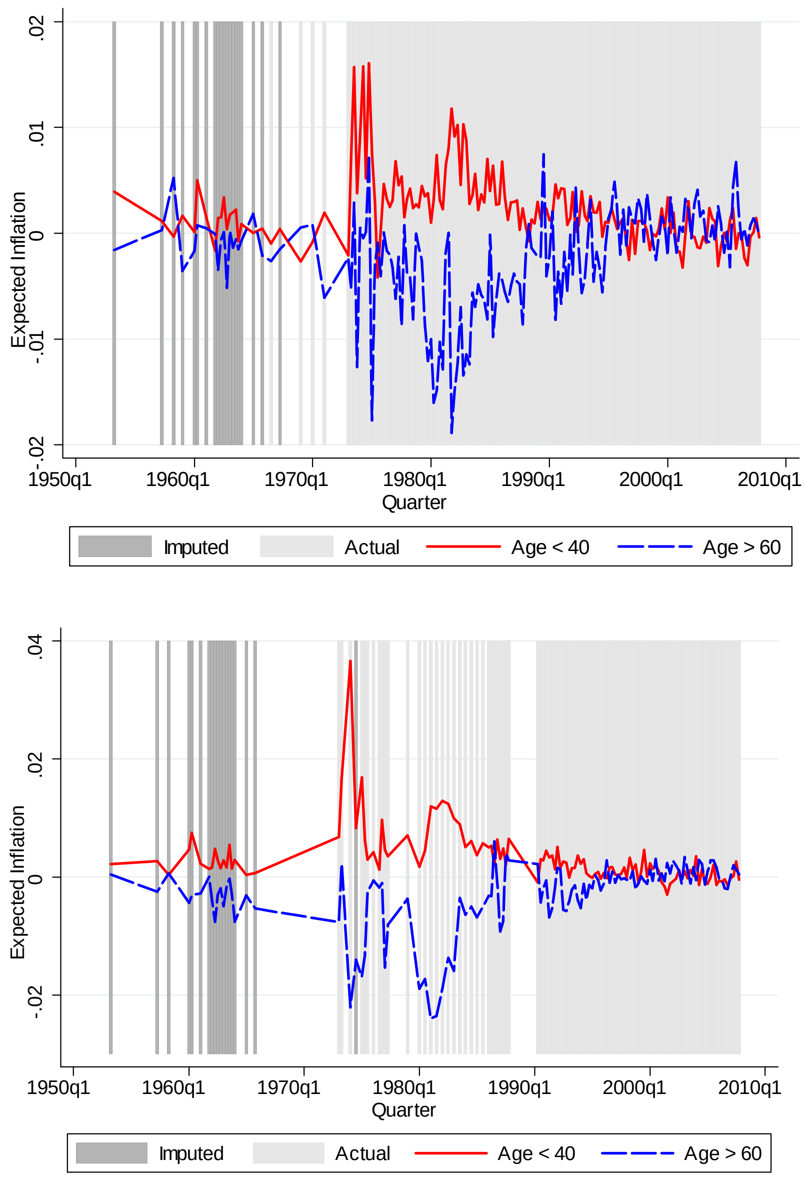

Figure 3: Average long-run inflation expectations of young and old relative to overall median

4 Estimated learning-from-experience effects

We now estimate the learning-from-experience effects by fitting Eq. (7) to the MCS inflation expectations data, aggregated at the cohort level. We measure $\tilde{\pi}^e_{t+1|t,s}$ with the inflation expectations reported by survey respondents of cohort s during quarter t. To make sure that we only use information in the explanatory variables that was available at the time the respondents made their forecasts, we calculate $\hat{\tau}_{t+1|t,s}$ with inflation data up and including quarter $t-1$. Effectively, we assume that individuals make their forecasts during the quarter based on information available at the beginning of the quarter. Note that we abuse notation here to some extent, as the forecasts in the survey data are over horizons of 1 year or 5-10 years, not one quarter, as in Eq. (7).

We estimate Eq. (7) with nonlinear least squares. To account for possible serial correlation of residuals within cohorts, we report standard errors that are robust to within-cohort clustering. Table 1 presents the estimation results for one-year expectations. Using the full sample, our estimate of the gain parameter in column (1) is $\theta = 2.33$ (s.e. 0.10) Comparing this estimate of θ with the earlier Figure 1, the weighting of past data is somewhere inbetween the weighting scheme for $\theta = 1$ (equal weights) and $\theta = 3$ (moderately fast declining weights). Roughly, the estimate implied weights that are declining a bit faster than linearly. The results in Table 1 also show that there is a strong relationship between $\hat{\tau}_{t+1|t,s}$ and inflation expectations, captured by the sensitivity parameter β, which we estimate to be 0.58 (s.e. 0.05). This magnitude of the β parameter implies that when two individuals differ in their weighted-average inflation experiences by 1%, their one-year inflation expectations tend to differ by 0.58% on average. Importantly, if there weren't any differences in expectations between individuals' with different inflation experiences (for example, because they all learned from the same historical data set in the same way, with the same forecasting rules) then β would be zero, because all the effect of historical inflation rates on current forecasts would be picked up by the time dummies. The fact that the point estimate of β is significantly different from zero is therefore direct evidence that historical data has different impact on individuals subjective expectations depending on whether the data was realized during their life-times or not.

Interestingly, the weighting of experienced inflation data implied by the point esti-

Table 1: Nonlinear least-squares estimates with 1-year inflation expectations, 1953-2007.

	(1) full sample	(2) w/o imputed data
Gain θ	2.33	2.32
	(0.10)	(0.11)
Sensitivity β	0.58	0.58
	(0.05)	(0.05)
Time dummies	Yes	Yes
Adj. R^2	77.8%	73.4%
#Obs.	7,765	7,200

Notes: Standard errors in parentheses are robust to heteroskedasticity and within-cohort clustering

mate of θ is remarkably similar to the weighting implied by the estimates obtained in Malmendier and Nagel (2008) from data on asset holdings. Explaining differences in allocation to inflation-sensitive long-term bonds vs. short-term cash-like instruments with experienced inflation also yields implied weights on past inflation similar to the ones we find here.[4]

To check whether the method used to impute percentage responses from categorical responses in periods when percentage responses are not available has any influence on the results, we also re-run the estimation without the imputed data. The results are presented in column (2). As can be seen, whether or not imputed data is used has little effect on the results.

Table 2 reports the same nonlinear least squares regressions, but now with 5-10 year inflation expectations as the dependent variable. The 5-10 year expectations data is available in fewer time periods, and so the number of observations in this table is about 2,000 lower than in Table 1. The gain parameter estimate in column (1) of $\theta = 2.15$; s.e. 0.10) is very close to the one we obtained with one-year expectations. The estimate $\beta = 0.76$ (s.e. 0.07) is higher than before, however, suggesting that long-run expectations may be more sensitive to learning-from-experience effects than one-year expectations. Again, whether we use the full sample or imputed data does not make much difference: The results in columns (1) and (2) are similar.

[4]The weighting function in Malmendier and Nagel (2008) is controlled by a parameter λ which relates to θ as $\theta \approx \lambda + 1$ (see Appendix A.3).

Table 2: Nonlinear least-squares estimates with 5-10 year inflation expectations, 1953-2007.

	(1)	(2)
	full sample	w/o imputed data
Gain θ	2.15	2.20
	(0.10)	(0.10)
Sensitivity β	0.76	0.74
	(0.07)	(0.07)
Time dummies	Yes	Yes
Adj. R^2	57.7%	57.8%
#Obs.	5,751	5,150

Notes: Standard errors in parentheses are robust to heteroskedasticity and within-cohort clustering.

The reported R^2 in Tables 1 and 2 include the effect of the time dummies, and so they are of limited use in judging the explanatory power of learning-from-experience. To get a better sense of how well the learning-from-experience effects explain cross-sectional differences in inflation expectations, 4 presents some plots of fitted values for different age groups.

For the purpose of these plots only, we average inflation expectations and the fitted values within the same age categories that we used earlier in Figure 3, i.e. within one category "young" (averaging across all cohorts that are in the age range from 25 to 39) and one "old" (averaging across cohorts that are in the age range 61 to 75). Since our estimation focuses on cross-sectional differences, due to the inclusion of time dummies, we also present the inflation expectations and fitted values for the "young" and "old" groups after subtracting the overall (interpolated) median each period. Thus, the plots also focus on cross-sectional differences, just like the estimation.

The top graph in Figure 4 plots the fitted values one-year expectations, i.e. those corresponding to the estimates in Table 1. Fitted values are drawn as dashed lines, raw inflation expectations as solid lines. The plot shows that the learning-from-experience model does a good job of explaining the differences between young and old in their inflation expectations. In particular, it accounts, to a large extent, for the large difference in opinion between young and old in the late 1970s and early 1980s.

The bottom graph in Figure 4 plots an equivalent graph using 5-10 year inflation expectations and the corresponding fitted values based on the estimates in Table 2.

Figure 4: Comparison of fitted 1-year (top) 5-10 year (bottom) inflation expectations for young and old with decreasing gain learning from life-time experience

Table 3: Controlling for age and age-specific inflation rates, 5-10 year inflation expectations, 1984-2007.

	1-year expectations		5-10 year expectations	
	(1)	(2)	(3)	(4)
Gain θ	2.18	4.16	1.86	3.12
	(0.13)	(0.34)	(0.10)	(0.38)
Sensitivity β	0.34	0.39	0.34	0.32
	(0.04)	(0.05)	(0.04)	(0.05)
(Age/1000)		-0.06		-0.07
		(0.02)		(0.02)
(Age/1000)$\times(\pi_{t-1}^{Elderly} - \pi_{t-1})$		-12.65		-3.74
		(4.68)		(5.64)
Time dummies	Yes	Yes	Yes	Yes
Adj. R^2	32.9%	33,4%	34.5%	34.9%
#Obs.	4,800	4,800	4,150	4,150

Notes: The sample runs from 1984:1 - 2007:4, the period for which lagged 12-month inflation rates from the experimental CPI for the elderly is available. Standard errors in parentheses are robust to heteroskedasticity and within-cohort clustering.

Here, too, the learning-from-experience model accounts for much of the difference in inflation expectations between young and old, although it is less successful than for one-year expectations in explaining some of the more extreme differences, such as the big spike in the mid-1970s.

One possible alternative story for these age-related differences in inflation expectations is that different age groups consume different consumption baskets, with different inflation rates, and that they form their expectations of future inflation based on the inflation rates they observe for their consumption baskets. If the inflation differentials between young and old individuals' consumption basket happen to be correlated with the differences in their experienced inflation rates, we could have an omitted variables problem. To address this issue, we control for differences between inflation rates on consumption baskets of the elderly and overall CPI inflation rates. We measure the inflation rates of the elderly from the experimental CPI for the elderly series (CPI-E) provided by the Bureau of Labor Statistics. We calculate inflation rates over 12-month periods at quarterly frequency, similar to our calculation of overall CPI inflation rates. We then include in our regressions the differential between the CPI-E and CPI inflation rates, $\pi_{t-1}^{Elderly} - \pi_{t-1}$, interacted with age.

Table 3 presents the results, with 1-year expectations in the first two columns, and 5-10 year expectations in columns (3) and (4). The 12-month inflation series based on the CPI-E is only available from the end of 1983 onwards, and so the sample in this table is restricted to 1984:1 to 2007:4. As a basis for comparison, we therefore first re-run the regression without the additional age-dependent inflation control on this shorter sample. The results in column (1) show that the estimate of the gain parameter is very similar to the earlier estimate in Table 1, but the sensitivity parameter β is estimated to be lower than before, but its magnitude is still statistically, as well as economically significant. In column (2) we then add the interaction term between age-related inflation differentials and age, as well as age itself (the $\pi_{t-1}^{Elderly} - \pi_{t-1}$ variable itself without the interaction is absorbed by the time dummies). First, the regression now attributes part of the differences in inflation expectations between young and old to a constant age effect via a negative coefficient on the age term. It is, however, not quite clear how to interpret this negative coefficient, since there doesn't seem to be a theory for why older individuals should have permanently lower inflation expectations than younger individuals. More importantly, we obtain a negative coefficient on the interaction term, which is not consistent with the idea that inflation expectations of the elderly may be positively related to the inflation rates on the consumption basket of the elderly. And while the coefficient is statistically significant from zero, the economic magnitude of the estimated effect is tiny. With an age differential of 50 years, the difference in coefficients on $\pi_{t-1}^{Elderly} - \pi_{t-1}$ would amount to $50/1000 = 0.05$, and so a 1% difference in $\pi_{t-1}^{Elderly} - \pi_{t-1}$ would translate into 0.05% difference in inflation expectations. Thus, economically, one should regard the estimate of the effect as virtually zero. Including age and the interaction term does, however, have some effect on the estimates for θ. With 4.16, the point estimate is substantially higher than in column (1), and the standard error is much higher as well.

Columns (3) and (4) repeat the same analysis for 5-10 year expectations, and the results are similar. The point estimate on the interaction term is again negative, but this time it is not only economically small, but also statistically not significantly different from zero. The increase of the point estimate of θ from including the interaction term is somewhat more moderate than with 1-year expectations. Overall, the evidence does not support the alternative theory that different consumption baskets could explain the heterogeneity in inflation expectations of young and old individuals.

5 Implications for forecast errors

We now investigate how learning-from-experience affects individuals' forecast errors. According to our estimation results, the young and old differ in how they are influenced by recent and distant inflation experiences. One would then expect that their success in forecasting inflation is different in different periods, depending on how realized inflation rates in those periods compare to the prior historical record.

5.1 Cross-section of forecast errors

Figure 5 compares the time-series of forecast errors of young and old by plotting forecast errors averaged within the group of individuals above 60 and below 40. The top graph shows forecast errors for one-year expectations, and the bottom graph those for 5-10 year expectations. For the latter, we use the average annual inflation rate realized over the subsequent 5 years to compute the forecast error, but the graph would look quite similar with a realized inflation rate calculated over a longer horizon, say, of 7 years.

Both graphs tell roughly the same story, but the differences between young and old are more pronounced for the longer term expectations in the bottom graph. In the 1970s, when inflation quickly shot up to levels not seen in decades, both young and old persistently underestimated future inflation for several years. However, older individuals underestimated it more so than younger individuals did. This is consistent with the learning-from-experience story in younger people put more weight on recent data, and therefore should be quicker to adapt to a dramatic and persistent change in inflation rates. Yet, in the 1980s, the high weight on relatively recent data of young people did not work in their favor. Older people, who had never adjusted their inflation expectations upwards as much as the young did, found themselves right on target with their forecasts as inflation quickly came back down in the early 1980s. Young people, instead, persistently overestimated future inflation for several years.

5.2 Aggregate forecast errors

Adaptive learning, and, in particular, learning-from-experience, may lead to predictable and persistent forecast errors (from the econometrician's perspective). If such forecast errors do not cancel out in the aggregate, they can influence macroeconomic outcomes.

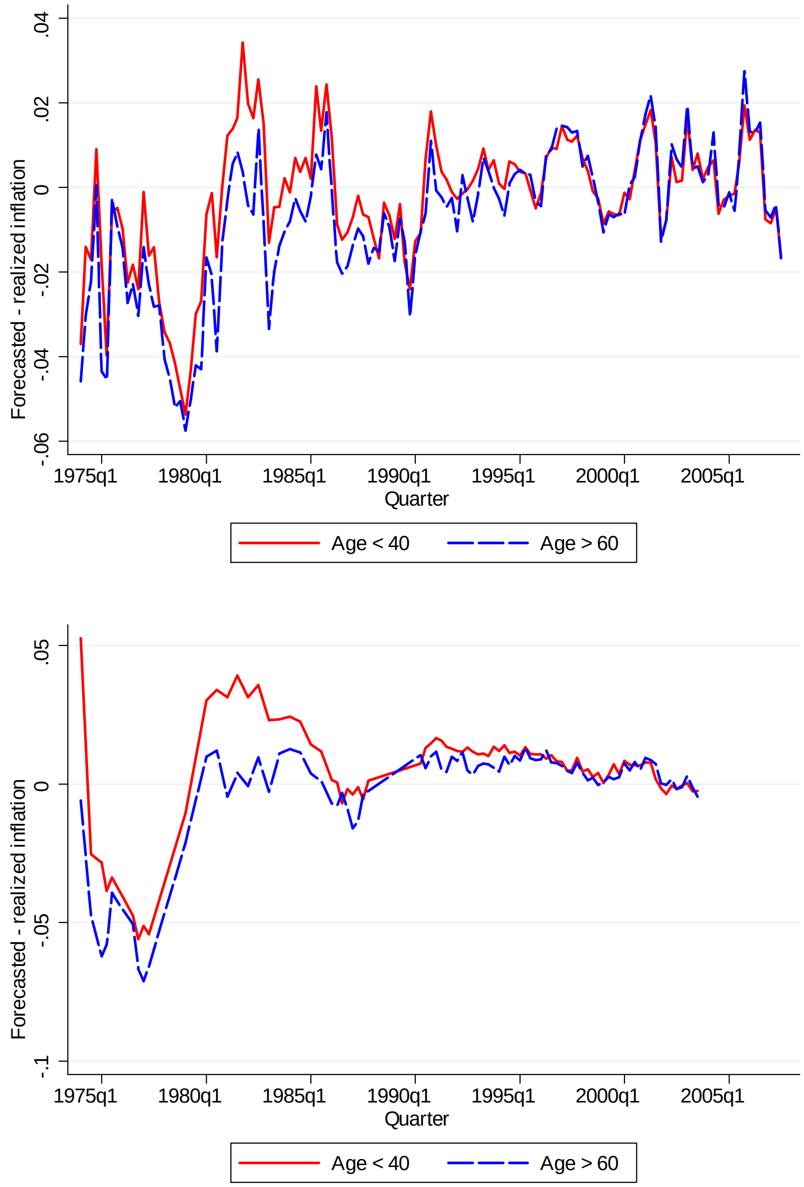

Figure 5: Forecast errors for one-year (top) and 5-10 year inflation expectations, defined as the forecast made in quarter t minus the subsequently realized average inflation rate over the forecast horizon

We therefore now turn our attention to the question whether we can link the learning-from-experience behavior to predictability in aggregate forecast errors. If aggregate inflation forecast errors are predictable, then one might also expect that returns of assets whose prices are sensitive to inflation expectations might have predictable returns, and so we also look at predictability of excess returns on long-term government bonds.

That learning-from-experience can lead to predictable and persistent forecast errors can be seen in the following simplified example. Suppose that an economic agent, born at time s, forecasts an *iid* random variable x_{t+1} that has unconditional mean μ. In the spirit of learning-from-experience, the agent uses, the sample mean of realizations $x_s, ..., x_t$ as the least-squares forecast, $x^e_{t+1|t}$, of x_{t+1}. In contrast, the econometrician who analyzes subjective expectations data ex-post, has data available extending beyond the $[s, t]$ interval. For simplicity, assume that the econometrician has a sufficiently large sample so that we can equate the sample mean observed by the econometrician with μ. In this case, the sample mean of the observed subjective forecasts $x^e_{t+1|t}$ also equals μ. Define the forecast errors as $x^e_{t+1|t} - x_{t+1}$. The econometrician then observes that $Cov(x^e_{t+1|t} - x_{t+1}, x^e_{t+1|t}) = Var(x^e_{t+1|t})$, which means that forecast errors are predictable by $x^e_{t+1|t}$. If individuals, on top of this, also suboptimally use the data realized during their life-times in constructing their forecasts, there can of course be additional sources of predictability.

For the purposes of econometrically identifying the learning-from-experience effect, we focused on cross-sectional differences in inflation expectations, and their evolution over time. However, given the estimates of how individuals weight their past experiences that we obtained from the estimation, we can now also ask what the aggregate implications of learning-from-experience might be. There are times when the average individual has high experienced inflation, and so the average $\hat{\tau}_{t+1|t,s}(\theta)$ is high in the population, and there are times when the average individual has low $\hat{\tau}_{t+1|t,s}(\theta)$. We look at whether this average, or aggregate, of $\hat{\tau}_{t+1|t,s}(\theta)$ is useful in predicting forecast errors and long-term bond returns.

We first set θ equal to 2.15, its point estimate from Table 2, column (1). To aggregate $\hat{\tau}_{t+1|t,s}(\theta)$ across individuals, we weight each age group by the proportion of aggregate liquid assets held by that age group. A liquid-assets-weighted average is appropriate as we ultimately want to look at the implications for asset prices, and households with more liquid assets presumably have a bigger impact on asset prices.

We obtain the liquid assets distribution from the Survey of Consumer Finances (SCF), as in Malmendier and Nagel (2008). Malmendier and Nagel (2008) find that the wealth distribution across age groups is quite stable, so we use their average wealth distribution from 1968 to 2004 in all periods to aggregate $\hat{\tau}_{t+1|t,s}(\theta)$ across age groups. The aggregated value is denoted by $\hat{\tau}_{t+1|t}(\theta)$ (i.e., omitting the s-subscript) which we refer to as *aggregate* experienced inflation. We aggregate subjective inflation expectations $\tilde{\pi}^e_{t+1|t,s}$ in the same way, and we denote them by $\tilde{\pi}^e_{t+1|t}$.

Dependent variable in our regressions is the forecast of inflation made during quarter t, aggregated across all age groups in the Michigan survey with age-group wealth weights from the SCF, minus the average annual inflation rate realized over the 12 months following the interview month in case of 1-year expectations, and five years in case of 5-10 year expectations. The regressions examine whether $\hat{\tau}_{t+1|t}(\theta)$, which is based on experienced inflation rates up to the end of quarter $t-1$, predicts these forecast errors. Since our imputation of percentage responses only targeted cross-sectional differences (our imputation regressions have time dummies; see Appendix A.2), but not the average level of percentage expectations, we omit all periods from these regressions in which we only have categorical inflation expectations data, and so we start the sample in 1973:1. To be able to use a Newey-West covariance matrix to correct for autocorrelation, we replace data in quarters in which survey information is missing with lagged (i.e., stale) values.

Table 4 presents the results. As column (1) shows, there is a strong positive relationship between aggregate experienced inflation and forecast errors. The coefficient estimate of 1.38 implies that a 1% higher weighted-average of individuals' experienced inflation is associated with a 1.38% higher forecast error in terms of the inflation rate in the year following the month in which the forecast was made. The R^2 of 20.5% shows that a substantial portion of the forecast error is predictable. On the other hand, $\hat{\tau}_{t+1|t}(\theta)$ is quite persistent this is reflected in the relatively high standard error of 0.52. In column (2), we add the annual inflation rate measured at the end of quarter $t-1$. It does not add explanatory power, and it is clear that $\hat{\tau}_{t+1|t}(\theta)$, which depends on inflation up to and including quarter $t-1$, does not simply pick up predictability contained in the most recent annual inflation rate. Column (3) adds the aggregate subjective inflation expectation measured during quarter t. Its inclusion has little impact on the coefficient on $\hat{\tau}_{t+1|t}(\theta)$.

Table 4: Relation between aggregate forecast errors and aggregate experienced inflation

	(1)	(2)	(3)
Intercept	-0.07	-0.05	-0.05
	(0.02)	(0.03)	(0.03)
Aggregate experienced inflation $\hat{\tau}_{t+1\|t}(\theta)$	1.38	1.17	1.11
	(0.52)	(0.52)	(0.52)
Lagged inflation π_{t-1}		-0.17	-0.10
		(0.08)	(0.12)
Aggregate inflation expectations $\tilde{\pi}^e_{t+1\|t}$			-0.16
			(0.27)
Adj. R^2	20.5%	28.2%	28.4%
#Obs.	137	137	137

Notes: OLS regressions with quarterly data from 1973:1 to 2007:2. Dependent variable is the forecast of 1-year inflation made during quarter $t+1$, aggregated across all age groups in the *MSC* with age-group wealth weights from the *SCF*, minus the average annual inflation rate realized over the 12 months following the interview month. Dependent and explanatory variables are left stale in quarters in which percentage expectations were imputed from categorical responses, or in which survey observations are unavailable. Newey-West standard errors in parentheses are robust to heteroskedasticity and autocorrelation (with 15 lags).

Table 5: Relation between aggregate forecast errors and aggregate experienced inflation

	(1)	(2)	(3)
Intercept	-0.09	-0.10	-0.14
	(0.04)	(0.05)	(0.04)
Aggregate experienced inflation $\hat{\tau}_{t+1\|t}(\theta)$	2.03	2.09	2.52
	(0.87)	(0.95)	(0.77)
Lagged inflation π_{t-1}		0.04	-0.26
		(0.10)	(0.11)
Aggregate inflation expectations $\tilde{\pi}^e_{t+1\|t}$			0.87
			(0.30)
Adj. R^2	32.0%	31.7%	49.6%
#Obs.	121	121	121

Notes: OLS regressions with quarterly data from 1973:1 to 2003:2. Dependent variable is the forecast of 5-10 year inflation made during quarter $t+1$, aggregated across all age groups in the *MSC* with age-group wealth weights from the *SCF*, minus the average annual inflation rate realized over the five years following the interview month. Dependent and explanatory variables are left stale in quarters in which percentage expectations were imputed from categorical responses, or in which survey observations are unavailable. Newey-West standard errors in parentheses are robust to heteroskedasticity and autocorrelation (with 15 lags).

Table 5 repeats the analysis for 5-10 year expectations. The sample now ends in quarter 2003:2 because we need at least five years ofrealized inflation data subsequent to the month in which the forecast was made. With 2.03 (s.e. 0.87), the coefficient estimate in column (1) is now moderately higher than with 1-year expectations data and the R^2 shows that a greater portion of the forecast error is predictable. As in Table 4, adding lagged inflation or aggregate inflation expectations has little effect on the coefficient on $\hat{\tau}_{t+1|t}(\theta)$. Overall, the data is consistent with the conjecture that individuals' tendency to extrapolate their forecast from life-time inflation experiences leads to predictable forecast errors.

6 Implications for bond returns

To see whether the predictability patterns in aggregate forecast errors translate into predictability in long-term bond returns, we run regressions similar to those in Tables 4 and 5, but now with excess returns on long-term bonds over the return on 1-month Treasury-Bills as the dependent variable.

Since the only data needed to generate $\hat{\tau}_{t+1|t}(\theta)$ are historical inflation rates and the liquid assets weights (assuming an estimate of θ has been obtained), there is no need to restrict the sample to the time period covered by the MSC, and we extend the sample back to returns spanning from 1926:1 (when our T-Bill return data becomes available) to 2007:4.

We obtain two series of long-term bond returns. One is a 10-year U.S. Treasury Bond total return index from Global Financial Data (GFD). For robustness, we also check the results with a return series of U.S. Treasury Bonds with maturities between 61 and 120 months from the Fama Bond database at the Center for Research in Security Prices ($CRSP$). The T-Bill return series is from Ibbotson Associates. We use quarterly returns as well as returns compounded to annual returns in our regressions.

The predictor variable is highly persistent, and its innovations (which are closely related to innovations in inflation) are likely to be correlated with long-term bond returns. Under these circumstances, it is well known that inference based on t-statistics relying on the usual first-order asymptotics leads to hypothesis tests that reject the null of no predictability too frequently in finite samples (Stambaugh (1999)). For this reason, we construct confidence intervals for the coefficient on the predictor variable

Table 6: Forecasting excess bond returns with aggregate experienced inflation

	1926:1 - 2007:4 (GFD)		1952:1-2007:4 (GFD)		1952:1-2007:4 (CRSP)	
	Qtrly. (1)	Ann. (2)	Qtrly (3)	Ann. (4)	Qtrly. (5)	Ann. (6)
OLS coeff. on $\hat{\tau}_{t+1\|t}(\theta)$	0.73	10.16	2.33	15.26	1.61	12.10
OLS s.e.	(0.64)	(2.67)	(1.16)	(4.74)	(0.85)	(3.41)
90% Bonferroni CI	[-0.31, 1.80]	[4.80, 13.72]	[0.65, 4.52]	[3.84, 19.52]	[0.36, 3.19]	[4.06, 15.38]
Adj. R^2	0.1%	14.4%	1.4%	14.8%	1.1%	17.7%
AR order of $\hat{\tau}_{t+1\|t}(\theta)$ by BIC	2	2	5	2	5	2
90% CI for largest AR root	[1.00, 1.00]	[1.00, 1.02]	[1.00, 1.01]	[1.00, 1.04]	[1.00, 1.01]	[1.00, 1.04]
#Obs.	328	82	224	56	224	56

Notes: Quarterly and annual regressions of U.S. Treasury bond returns in excess of 1-month Treasury Bill returns on beginning of period values aggregate experienced inflation. Quarterly and annual bond returns are calculated by compounding the monthly returns. The 1974:1-2007:4 sample excludes quarters in which percentage expectations were imputed from categorical responses. The table shows OLS estimates along with a 90% Bonferroni confidence interval following Campbell and Yogo (2006) for the coefficient on aggregate experienced inflation.

using the methods of Campbell and Yogo (2006). Campbell and Yogo use local-to-unity asymptotics to achieve a better approximation of the finite-sample distribution when the predictor variable is persistent. Their construction of the confidence interval uses the Bonferroni method to combine a confidence interval for the largest autoregressive root of the predictor variable with confidence intervals for the predictive coeffficient conditional on the largest autoregressive root.

The results are presented in Table 6. Columns (1) and (2) show the estimated coefficients and 90% confidence intervals with quarterly and annual data in the long sample. While the evidence for predictability is somewhat weak with quarterly data (the confidence interval from -0.31 to 1.80 includes 0.00), the evidence is much stronger with annual data, with the confidence interval for the predictive coefficient stretching from 4.80 to 13.72. The annual regressions produce a predictive adjusted R^2 of 14.4%, which shows that a considerable portion of bond excess return is predictable with aggregate inflation experiences. When individuals on average have experienced high inflation during their life-times, expected excess returns are higher. This is consistent with the earlier evidence that high inflation experiences lead to higher inflation expectations, and overestimation of future inflation rates, which depresses prices of long-term bonds and raises their expected excess returns.

To check the reliability of the data series on total returns of long-term bonds, we also run the predictive regressions with the $CRSP$ long-term bond returns, which is available for the 1952:1 to 2007:4 time period. To compare the two data series, we also first run regressions with the GFD data for this shorter sample, with the results shown in columns (3) and (4). The results are stronger than for the longer sample. Both with quarterly and annual data, the confidence interval for the predictive coefficient now excludes zero, and the adjusted R^2 is a little higher, particularly with quarterly data. Switching to the $CRSP$ data in columns (5) and (6) produces almost identical results. Overall, the evidence supports the hypothesis that aggregate inflation experiences predict bond excess returns.

At the bottom of the table, we also report the estimated autoregressive lag length for the predictor variable, as determined by the Bayesian Information Criterion (BIC), as well as a confidence interval for its largest autoregressive root. These are among the inputs to Campbell and Yogo's construction of confidence intervals. The estimated lag length is greater than one in each case, which means that would not be appropriate to assume an AR(1) process for the predictor, as in Stambaugh (1999), for example. The confidence intervals for the largest autoregressive root contain an explosive root in each case. This is similar to the dividend-price ratio, for example (see Campbell and Yogo (2006)), which underscores the importance of accounting for the persistence of the predictor variable in testing for predictability.

To check the robustness of the results, we have also re-run the predictive regressions with a different timing convention for the predictor variable $\hat{\tau}_{t+1|t}(\theta)$. So far, we computed $\hat{\tau}_{t+1|t}(\theta)$ with inflation rates up to the end of quarter $t-1$, assuming that this is the information that is available to individuals making forecasts in quarter t. We get similar results if we use inflation rates up to the end of quarter t, which assumes that individuals can guess the inflation rate before it is officially released, perhaps from their own consumption experiences. This issue has some relevance because the Campbell-Yogo method assumes that innovations in the predictor and the return may be contemporaneously correlated, but are iid, i.e. there is no (cross-) serial correlation.

A priori it is not quite clear which of the two timing conventions best conforms to this assumption, so it is good news that it works both ways. Also, for annual data this timing issue should be less of a problem.

7 Discussion and conclusion

Our empirical analysis shows that individuals' inflation expectations have a component that depends on past life-time inflation experiences. This component can be interpreted as the result of adaptive learning from life-time experiences. It has an important influence on the dynamics of subjective expectations, and we find that it has some predictive power for individuals' aggregate forecast errors and the excess returns on long-term bonds. We conclude by discussing some connections between our findings and some open issues in macroeconomics and finance.

The learning-from-experience mechanism can produce heterogeneity in expectations. In particular, it can help explain the increased disagreement between individuals during the high-inflation periods of the 1970s and early 1980s noted in Mankiw and Wolfers (2003) as disagreement between young and old individuals, who relied to different extent on extrapolation of recent inflation data in forming their expectations. It thus provides an alternative to the "sticky information" hypothesis in Mankiw and Reis (2002) and Carroll (2003) in explaining this heterogeneity.

Moreover, just like "sticky information", adaptive learning provides a possible explanation for inertia in inflation expectations, which in turn can help understand persistence inflation rates (see, e.g., Roberts (1997); Orphanides and Williams (2006); Milani (2007)). Unlike standard adaptive learning rules, however, learning-from-experience implies that there is cross-sectional heterogeneity in forecasts. This heterogeneity allows us to estimate the forecasting rule while at the same time controlling for other common macroeconomic factors that might also influence expectations.

The learning-from-experience results also suggests a way in which economic agents' memory is bounded or decaying: The memory of macroeconomic history is lost as new generations emerge whose subjective beliefs are shaped by relatively recent experience. As a consequence, learning dynamics may be perpetual, without convergence in the long-run.

Finally, the fact that individuals' beliefs are heterogeneous due to their different life-time experiences of inflation suggests that they might trade in asset markets in a way consistent with those beliefs. Indeed, Malmendier and Nagel (2008) find that differences in macroeconomic experiences help explain differences in portfolio allocations to risky assets, and Piazzesi and Schneider (2008a) show that young and old individuals' asset allocation decision during the 1970s were consistent with their disparate views about future inflation. Our evidence that life-time experiences of inflation aggregated across individuals helps forecast excess returns of long-term bonds lends further support to the idea that the belief variation we find in the expectations data also affects asset prices.

References

Barsky, R. B. (1987): "The Fisher Hypothesis and the Forecastability and Persistence of Inflation," *Journal of Monetary Economics*, 19, 3–24.

Bernanke, B. S. (2007): "Inflation Expectations and Inflation Forecasting," Speech at NBER Monetary Economics Workshop, available at http://www.federalreserve.gov/newsevents/speech/Bernanke20070710a.htm.

Branch, W. A. (2004): "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations," *Economic Journal*, 114, 592–621.

——— (2007): "Sticky Information and Model Uncertainty in Survey Data on Inflation Expectations," *Journal of Economic Dynamics and Control*, 31, 245–276.

Bray, M. (1982): "Learning, Estimation, and the Stability of Rational Expectations Equilibria," *Journal of Economic Theory*, 26, 318–339.

Campbell, J. Y., and M. Yogo (2006): "Efficient Tests of Stock Return Predictability," *Journal of Financial Economics*, 81, 27–60.

Carroll, C. D. (2003): "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics*, 118(1), 269–298.

Cukierman, A., and P. Wachtel (1979): “Differential Inflationary Expectations and the Variability of the Rate of Inflation: Theory and Evidence,” *American Economic Review*, 69(4), 595–609.

Curtin, R. T. (1982): “Indicators of Consumer Behavior: The University of Michigan Surveys of Consumers,” *Public Opinion Quarterly*, 46(3), 340–352.

Evans, G. W., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.

Greenwood, R., and S. Nagel (2009): “Inexperienced Investors and Bubbles,” *Journal of Financial Economics*, forthcoming.

Honkapohja, S., and K. Mitra (2003): “Learning with Bounded Memory in Stochastic Models,” *Journal of Economic Dynamics and Control*, 27, 1437–1457.

Kalman, R. E. (1960): “A New Approach to Linear Filtering and Prediction Problems,” *Journal of Basic Engineering*, 82(1), 35–45.

Kaustia, M., and S. Knüpfer (2008): “Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions,” *Journal of Finance*, 63, 2679–2702.

Keynes, J. M. (1936): *The General Theory of Employment, Interest, and Money*. Macmillan, London.

Malmendier, U., and S. Nagel (2008): “Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?,” working paper, Stanford University and UC Berkeley.

Mankiw, G. N., and R. Reis (2002): “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 117(4), 1295–1328.

Mankiw, N. Gregory, R. R., and J. Wolfers (2003): “Disagreement About Inflation Expectations,” in *NBER Macroeconomics Annual 2003*, ed. by M. Gertler, and K. Rogoff.

Marcet, A., and J. Sargent, Thomas (1989): "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–368.

Milani, F. (2007): "Expectations, Learning, and Macroeconomic Persistence," *Journal of Monetary Economics*, 54, 2065–2082.

Muth, J. F. (1960): "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association*, 55(290), 299–306.

Nelson, C. R., and G. W. Schwert (1977): "Short-Term Interest Rates as Predictors of Inflation: On Testing the Hypothesis that the Real Rate of Interest Is constant," *American Economic Review*, 67, 478–486.

Orphanides, A., and J. C. Williams (2006): "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations," *Journal of Economic Dynamics and Control*, 29, 1927–1950.

Piazzesi, M., and M. Schneider (2008a): "Inflation and the Price of Real Assets," working paper, Stanford University.

——— (2008b): "Trend and Cycle in Bond Premia," working paper, Stanford University.

Roberts, J. M. (1997): "Is Inflation Sticky?," *Journal of Monetary Economics*, 39, 173–196.

Sargent, T. J. (1993): *Bounded Rationality in Macroeconomics*. Clarendon Press, Oxford, UK.

Shiller, R. J. (2005): *Irrational Exuberance*. Princeton University Press, Princeton, NJ, second edn.

Stambaugh, R. F. (1999): "Predictive Regressions," *Journal of Financial Economics*, 54, 375–421.

Stock, J. H., and M. W. Watson (2003): "Why Has U.S. Inflation Become Harder to Forecast?," *Journal of Money, Credit and Banking*, 39(1), 3–33.

Vissing-Jorgensen, A. (2003): “Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions,” in *NBER Macroeconomics Annual*.

APPENDIX

A.1 Michigan Survey data

The inflation expectations data is derived from the responses to two questions, the first is categorical, while the second one elicits a percentage response. For example, for 1-year expectations the two questions are:

1. "During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are right now?"
2. "By about what percent do you expect prices to go (up/down) on average during the next 12 months?"

As outlined in Curtin (1996), some adjustments to the raw data are necessary to address some known deficiencies. We follow Curtin's approach, which is also the approach used by the Michigan Survey in constructing its indices from the survey data:

For respondents who provided a categorical response of "up" ("down"), but not a percentage response, we drew a percentage response from the empirical distribution of percentage responses of those who gave the same categorical response of "up" ("down") in the same survey period. Prior to the February 1980 survey, respondents were not asked about percentage expectations if they responded (in the categorical first part of the question) that they expected prices to decline. We assign a value of -3% to these cases before February 1980. In most survey periods, they account for less than 2% of observations.

Starting in March 1982 the administrators of the Michigan survey implemented additional probing, which revealed that the categorical response that prices will remain the "same" was often misunderstood as the meaning that the inflation rate stays the same. We use the adjustment factors developed in Curtin (1996) to adjust a portion of "same" responses prior to March 1982 to "up", and we assign a percentage response by drawing from the empirical distribution of those observations in the same survey period with a categorical response of "up".

A.2 Imputation of percentage expectations from categorical responses

In the early years of the Michigan survey, only categorical responses about prices going "up", "down", or stay the "same" were elicited, but no percentage responses. We nevertheless attempt to use the information in those surveys in our analysis of percentage expectations by imputing percentage responses from the categorical information. We do so by estimating the relationship between categorical responses, the dispersion of categorical responses, and percentage responses in those periods in which we have both categorical and percentage response data. We conjecture that the average percentage response of individuals in an age group should be positively related to the proportion of "up" responses and negatively to the proportion of "down" responses.

We first calculate the proportion of "up" and "down" responses, $p^{up}_{s,t}$ and $p^{down}_{s,t}$,within each age group at each point in time. We then run a pooled regression of percentage inflation expectations, $\hat{\pi}^e_t$, on $p^{up}_{s,t}$ and $p^{down}_{s,t}$, including a full set of time dummies, and obtain, for one-year expectations, the fitted values

$$\hat{\pi}^{e,imp}_t = \text{...time dummies...} + \underset{(0.002)}{0.059} p^{up}_{s,t} - \underset{(0.004)}{0.019} p^{down}_{s,t} \quad (R^2 = 78.7\%)$$

and for 5-10 year expectations,

$$\hat{\pi}^{e,imp}_t = \text{...time dummies...} + \underset{(0.004)}{0.055} p^{up}_{s,t} - \underset{(0.007)}{0.036} p^{down}_{s,t} \quad (R^2 = 67.2\%)$$

with standard errors in parentheses that are robust for heteroskedasticity and within-cohort clustering.

Because we employ time dummies in our main analysis, our main concern here is whether the imputed expectations track well cross-sectional differences of expectations across age groups, rather than the overall mean over time, and so we also estimate the relationship between percentage expectations and categorical responses with time dummies included in the regression.

Figure A.1 illustrates how the imputed percentage expectations compare with the actual expectations in the time periods in which we have both categorical and percentage expectations data. To focus on cross-sectional differences between age groups, the figure shows the average fitted and actual values for individuals below 40 and above 60 years of age after subtracting the overall median of each time period.

A.3 Adaptive learning schemes and the weighting function of Malmendier and Nagel (2008)

Consider an individual of age $t-s$ making an inflation forecast at time t. The weighting function in Malmendier and Nagel (2008), implies that this individual forms a weighted average of past inflation, where the inflation rate observed at time $t-k$ (with $k \leq t-s$) gets the following weight:

$$\omega_{t-s}(k) = \frac{\left(\frac{t-s-k}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s}\left(\frac{t-s-j}{t-s}\right)^{\lambda}}. \tag{A.1}$$

This implies that the most recent observation, i.e. time-t inflation, π_t, receives the weight

$$\omega_{t-s}(0) = \frac{1}{\sum_{j=0}^{s}\left(\frac{t-s-j}{t-s}\right)^{\lambda}}. \tag{A.2}$$

For comparison, in a stochastic recursive learning algorithm with a gain sequence γ_t, we would have

$$\pi_t^e = (1-\gamma_t)\,\pi_{t-1}^e + \gamma_t \pi_t, \tag{A.3}$$

which implies that the most recent observation carries the weight $\tilde{\omega}_t(0) = \gamma_t$. In our case, the gain parameter is age-dependent instead of time-dependent, and so we write γ_{t-s}, so that a member of cohort s produces the forecast

$$\pi_{s,t}^e = \left(1-\gamma_{t-s}\right)\pi_{t-1}^e + \gamma_{t-s}\pi_t. \tag{A.4}$$

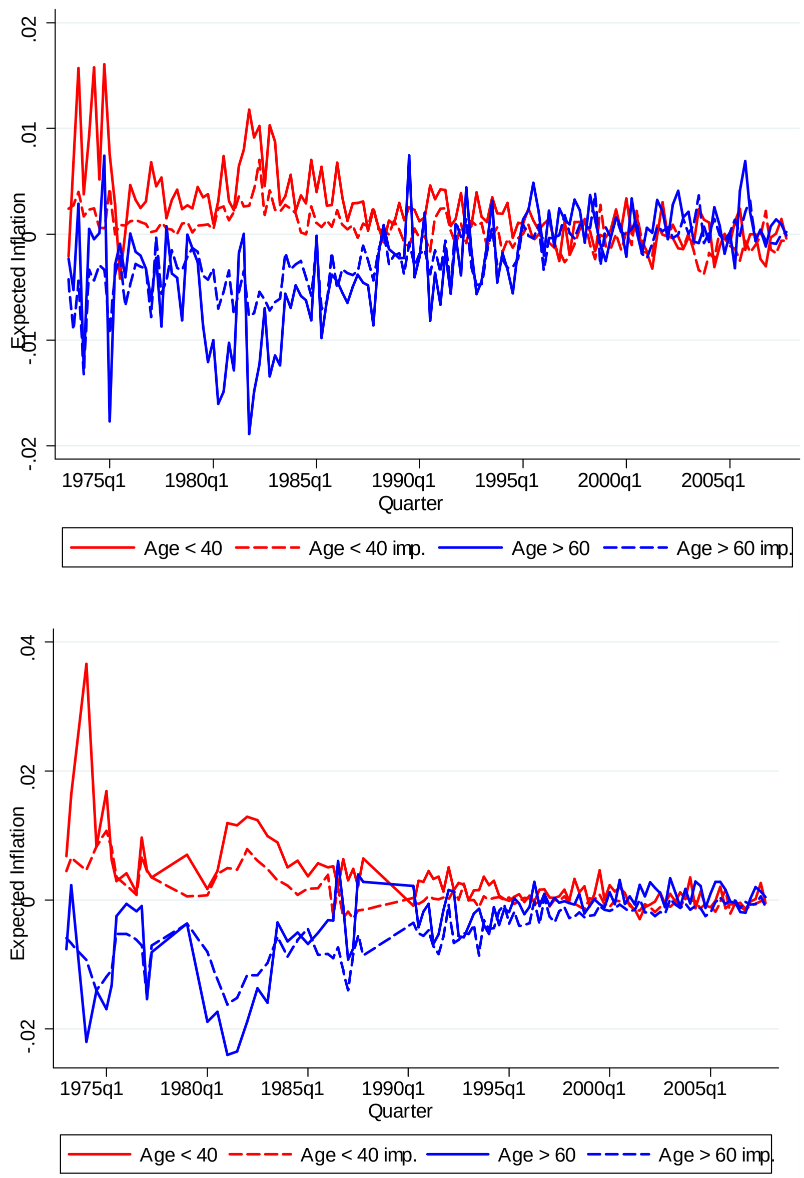

Figure A.1: Actual and imputed long-run percentage inflation expectations relative to their medians for 1-year expectations (top) and 5-10 year expectations (bottom)

Iterating, one finds that earlier observations receive the weight

$$\tilde{\omega}_{t-s}(k) = \begin{cases} \gamma_{t-s} \quad \text{for } k = 0 \\ \gamma_{t-s-k} \prod_{j=0}^{k-1} \left(1 - \gamma_{t-s-k}\right) \quad \text{for } k > 0 \end{cases} . \tag{A.5}$$

We now show that both weighting schemes are equivalent if the gain sequence is chosen to be age-dependent in the following way:

$$\gamma_{t-s} = \frac{1}{\sum_{j=0}^{t-s} \left(\frac{t-s-j}{t-s}\right)^{\lambda}} \tag{A.6}$$

We present a proof by induction. First, the choice of γ_{t-s} in Eq. (A.6) implies that $\tilde{\omega}_{t-s}(0) = \omega_{t-s}(0)$. It remains to be shown that if $\tilde{\omega}_{t-s}(k) = \omega_{t-s}(k)$, then $\tilde{\omega}_{t-s}(k+1) = \omega_{t-s}(k+1)$ (with $k \leq t-s$). Thus, assume that

$$\tilde{\omega}_{t-s}(k) = \frac{\left(\frac{t-s-k}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s} \left(\frac{t-s-j}{t-s}\right)^{\lambda}}. \tag{A.7}$$

Then, from Eq. (A.5),

$$\begin{aligned}
\tilde{\omega}_{t-s}(k+1) &= \gamma_{t-s-k-1} \frac{\left(1 - \gamma_{t-s-k}\right)}{\gamma_{t-s-k}} \tilde{\omega}_{t-s}(k) \\
&= \frac{\left[\sum_{j=0}^{t-s-k} \left(\frac{t-s-j}{t-s-k}\right)^{\lambda}\right] - 1}{\sum_{j=0}^{t-s-k-1} \left(\frac{t-s-j}{t-s-k-1}\right)^{\lambda}} \frac{\left(\frac{t-s-k}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s} \left(\frac{t-s-j}{t-s}\right)^{\lambda}} \\
&= \frac{\left[\sum_{j=0}^{t-s-k} \left(\frac{t-s-j}{t-s-k}\right)^{\lambda}\right] - 1}{\sum_{j=0}^{t-s-k-1} \left(\frac{t-s-j}{t-s-k}\right)^{\lambda}} \frac{\left(\frac{t-s-k-1}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s} \left(\frac{t-s-j}{t-s}\right)^{\lambda}} \\
&= \frac{\sum_{j=0}^{t-s-k-1} \left(\frac{t-s-j}{t-s-k}\right)^{\lambda}}{\sum_{j=0}^{t-s-k-1} \left(\frac{t-s-j}{t-s-k}\right)^{\lambda}} \frac{\left(\frac{t-s-k-1}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s} \left(\frac{t-s-j}{t-s}\right)^{\lambda}} \\
&= \frac{\left(\frac{t-s-k-1}{t-s}\right)^{\lambda}}{\sum_{j=0}^{t-s} \left(\frac{t-s-j}{t-s}\right)^{\lambda}} \\
&= \omega_{t-s}(k+1),
\end{aligned}$$

where for the third-to-last equality we multiplied numerator and denominator by $\left(\frac{t-s-k-1}{t-s-k}\right)^{\lambda}$. This concludes the proof.

Finally, we show that the gain sequence (A.6) can be approximated as

$$\gamma_{t-s} \approx \frac{\lambda+1}{t-s}.$$

To see this write the gain in Eq. (A.6) as

$$\gamma_{t-s} = \frac{(t-s)^{\lambda}}{\sum_{j=0}^{t-s}(t-s-j)^{\lambda}}.$$

Focusing on the denominator of this expression, note that if one were to make the increments j infinitesimally small (instead of being discrete steps of 1), the denominator would become $\int_0^{t-s} x^{\lambda}dx = \frac{1}{\lambda+1}(t-s)^{\lambda+1}$. Therefore, in this limiting case of infinitesimal increments, we get

$$\gamma_{t-s} = \frac{\lambda+1}{t-s}.$$

In our case it turns out that with quarterly increments, and $t-s$ at least 100 quarters, this approximation is, for all practical purposes virtually identical with the true gain sequence in Eq. (A.6).